


NVA Mapping Distribution Mechanism

draft-dunbar-nvo3-nva-mapping-distribution-00

Linda Dunbar

Donald Eastlake

Various ways of NVAs connected to NVEs



Locations:

- Embedded in routers/switches in the core, or as standalone servers attached to them.
- Standalone servers or VMs connected to Edges via the client side port

Contents:

- Centralized NVA
- Distributed NVA:
 - Each NVA has mapping for a subset of VNs
 - multiple NVAs have mapping entries for a VN

Push Model

- **Incremental Push Service Update**
 - Achieved by Link State Update to distribute the incremental updates.
- **Requesting Push Service:**
 - Push NVAs use VN scoped reliable Link State flooding to announce their availability to push mapping information.
 - NVEs use VN scoped reliable Link State flooding to announce all the Virtual Networks in which they are participating
 - Whenever, there are changes in the mapping entries, NVA uses CSNP messages to only send the changed portion of the entries.
- **Policies:** When ingress edge can't find entries for the incoming data frame:
 - simply drop the data frame,
 - flood it to all other edges that are in the same VN, or
 - start the “pull” process to get information from Pull NVA

Reachable Interface Addresses (IA) TLV

- To advertise a set of addresses within a VN being attached to (or reachable by) a specific NVE
- These addresses can be in different address families. For example, it can be used to declare that a particular interface with specified IPv4, IPv6, and 48-bit MAC addresses in some particular VN is reachable from a particular NVE.

```

+-----+
| Type = TBD                               | (2 bytes)
+-----+
| Length                                   | (2 bytes)
+-----+
| Addr Sets End                           | (2 bytes)
+-----+
| NVE Address subTLV ...                   | (variable)
+-----+
| Flags                                   | (1 byte)
+-----+
| Confidence                               | (1 byte)
+-----+
| Template ...                             | (variable)
+-----+
| Address Set 1 (size determined by Template) |
+-----+
| Address Set 2 (size determined by Template) |
+-----+
| ...
+-----+
| Address Set N (size determined by Template) |
+-----+
| optional sub-sub-TLVs ...
+-----+

```

Requesting Push Service

- When a NVE is initialized or re-started, it uses Virtual Network scoped instances of the IS-IS to announce all the Virtual Networks in which it is participating




Figure 2. Enabled-VN TLV

Incremental Push service

- A new TLV is needed for to carry CSNP timeout value and a flag for NVA to indicate it has completed all updates.





Figure 3. CSNP Complete TLV

Flags: A byte of flags defined as follows:



Pull Query Format

- **Pull Request** Triggered by:
 - An NVE receives an ingress data frame with a destination whose egress NVE is unknown, or
 - An NVE receives an ingress ARP/ND request for a target whose link address (MAC) or egress edge NVE is unknown.

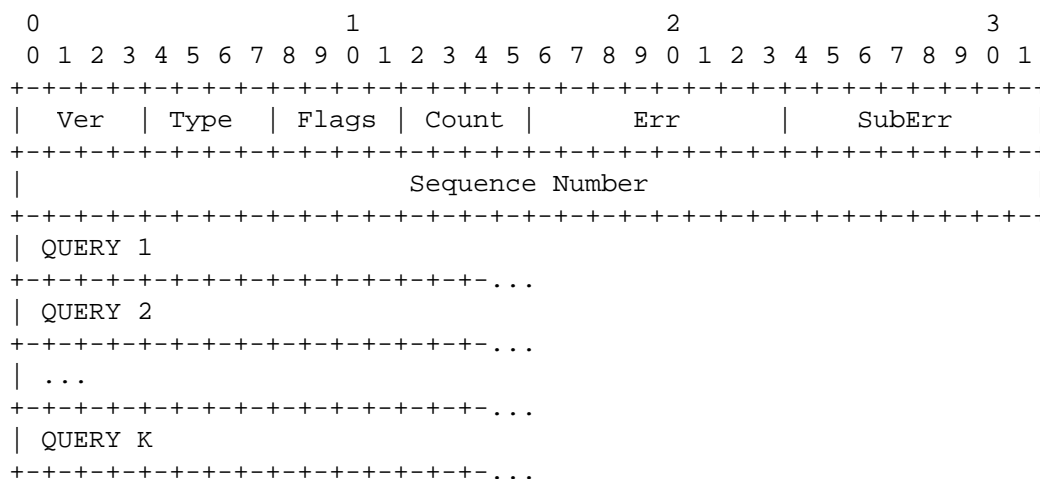


Figure 4. Pull Query TLV

Query Record

QUERY: Each QUERY Record within a Pull Directory Query message is formatted as follows:

```

  0  1  2  3  4  5  6  7  8  9 10 11 12 13 14 15
+-----+-----+-----+-----+-----+-----+-----+-----+
|           SIZE           |     RESV     |     QTYPE     |
+-----+-----+-----+-----+-----+-----+
| If QTYPE = 1
+-----+-----+-----+-----+-----+-----+
|                               AFN                               |
+-----+-----+-----+-----+-----+-----+
| Query address ...
+-----+-----+-----+-----+-----+-----+
| If QTYPE = 2, 3, 4, or 5
+-----+-----+-----+-----+-----+-----+
| Query frame ...
+-----+-----+-----+-----+-----+-----+

```

QTYPE	Description
0	reserved
1	address query
2	ARP query frame
3	ND query frame
4	RARP query frame
5	Unknown unicast MAC query frame
6-14	assignable by IETF Review
15	reserved

AFN: Address Family Number of the query address.

Pull Response

RESPONSE: Each RESPONSE record within a Pull NVA Response message is formatted as follows:

```
  0  1  2  3  4  5  6  7  8  9 10 11 12 13 14 15
+-----+-----+-----+-----+-----+-----+-----+-----+
|          SIZE          |OV| RESV |   Index   |
+-----+-----+-----+-----+-----+-----+-----+
|                   Lifetime                   |
+-----+-----+-----+-----+-----+-----+
|                   Response Data ...          |
+-----+-----+-----+-----+-----+-----+...
```

Push-Pull Hybrid Model

- Push model are used for some VNs, and pull model are used for other VNs.
 - It can be operator's decision (i.e. by configuration) on which VNs' mapping entries are pushed down from NVA (e.g. frequently used) and which VNs' mapping entries are pulled (e.g. rarely used).
 - Useful for Gateway nodes where great number of VNs are enabled.
- Or, a portion of hosts in a VN is pushed, other portion has to be pulled.